N90-25547

AUTOMATED CONTROL OF HIERARCHICAL SYSTEMS USING VALUE-DRIVEN METHODS

Dr. George E. Pugh
Chairman of the Board
Decision-Science Applications, Inc.
1110 North Glebe Road, Suite 400
Arlington, VA 22201-4714
(703) 243-2500

Thomas E. Burke
Director, Command and Control Applications
Decision-Science Applications, Inc.
1110 North Glebe Road, Suite 400
Arlington, VA 22201-4714
(703) 243-2500

ABSTRACT

This paper provides an introduction to the Value-driven methodology, which has been successfully applied to solve a variety of difficult decision, control, and optimization problems. Many real-world decision processes (e.g., those encountered in scheduling, allocation, and command and control) involve a hierarchy of complex planning considerations. For such problems it is virtually impossible to define a fixed set of rules that will operate satisfactorily over the full range of probable contingencies. DSA's value-driven methodology offers a systematic way of automating the intuitive, common-sense approach used by human planners. The inherent responsiveness of value-driven systems to user-controlled priorities makes them particularly suitable for semi-automated applications in which the user must remain in command of the systems operation. Three examples of the practical application of the approach in the automation of hierarchical decision processes are discussed: the TAC Brawler air-to-air combat simulation is a four-level computerized hierarchy; the autonomous underwater vehicle mission planning system is a three-level control system; and the Space Station Freedom electrical power control and scheduling system is designed as a two-level hierarchy. The methodology is compared with rule-based systems and with other more widely-known optimization techniques.

INTRODUCTION

The value-driven methodology described in this paper had its beginnings in the 1960s as an optimization technique for large non-linear, discrete problems; specifically for the allocation of strategic weapons to targets. For example, the assignment of U.S. missiles and bombers to Soviet targets has on the order of 400 weapons types, 30,000 targets, a non-linear objective function, and is defined over a discrete space.

Over the past two decades, it has been successfully applied to an exceptionally wide variety of very complex automation and control problems, including:

- Decision and control systems for autonomous robotic vehicles
- Network design and optimization for telecommunications systems
- Automated and semi-automated control systems for military command, control, and communications applications
- Production scheduling and optimization systems
- Scheduling and control systems for space applications
- Simulations of human decision processes in computerized combat simulations

DSA's value-driven control methodology focuses on the modeling of the system being represented rather than on the optimization technique itself. Furthermore, many problems cannot be represented as a strict optimization problem either because they are too complex or because optimization is not the objective but rather the desire is to have an automated or semi-automated system that performs the same as an *experienced* user would perform given the same set of conditions.

Over the years of research and application of the methodology, DSA's system development efforts have fallen into four major areas:

- pure optimization systems
- scheduling systems
- control systems--centralized, distributed, and hierarchical
- systems that mimic human decision process

This paper provides an introduction to the concepts underlying value-driven control methods and briefly describes three applications in the automated control of hierarchical systems. We also provide a brief comparison with other more widely known approaches for building intelligent systems.

PRINCIPLES OF VALUE-DRIVEN CONTROL

A value-driven control system is one in which the automated decision processes are governed by value-maximizing principles rather than by the "if-then" rules common to expert systems and traditional automation software. Value-driven systems are designed to maximize a system of priorities, and are automatically and intelligently responsive to real-time changes in the problem environment The systematic value-maximizing process which is used in place of a pre-defined set of decision rules operates essentially as follows:

1. It "considers," or searches over, a set of possible decision alternatives;

2. It evaluates each alternative in terms of the currently specified value priorities, and

3. It selects and implements the alternative that yields the maximum value in terms of the currently specified value priorities.

Use of Decision Entities

Decomposition of a complex problem is an essential step in the formulation of a value-driven system.

Indeed, value-driven systems that are responsible for complex decision and control processes are typically composed of a large number of automated "decision entities", each responsible for the functions within a limited subarea. System behavior is then controlled by a system of value functions, one value function per decision entity. A typical value function will have the functional form

$$V = V(a_i, \ldots a_n, x_i, \ldots x_n) \quad (1)$$

where the a_i's are adjustable parameters controlled by the user or by higher levels in a hierarchical system and the x_i's are the state variables of a projected state corresponding to one possible decision alternative.

Each decision entity will have a suitable representation of the decision environment, so that it can formulate feasible decision alternatives and evaluate the desirability of the alternatives in terms of the currently applicable objectives and priorities. A decision entity will select an alternative that maximizes the value function. In particular, the responsible decision entity must be provided with:

- Appropriate policy guidance concerning the current objectives and priorities, and
- Current situation or status information concerning the assigned realm of decision responsibility.

Goal Orientation Through Modifiable Subgoals

The a_i's of (1) are the means by which the set of priorities and objectives that guide each decision entity are adjusted to reflect the goals of the system. Thus, each decision entity is operating in the realm of subgoals that are modified as the current state of the system changes from decision point to decision point.

The formulation of the priority scheme within a value-driven system and the refinement of the values are essential to the correct solution to a problem. The central objective is to ensure that the structure of values will in fact be effective in achieving the real goal for the system.

Responsiveness to Command Priorities

The inherent ability of value-driven systems to respond intelligently to command priorities makes it possible to develop hierarchical control systems in which each level of an automated hierarchy is responsive to the changing objectives and priorities specified at higher levels. A high level entity, by changing the a_i's in the value function of a lower level entity, changes the lower level entity's perception of the relative desirability of different goals or objectives, and thereby influences the behavior of the lower level's selection of courses of action.

The responsiveness of value-driven systems to user-controlled priorities makes such systems particularly suitable for semi-automated applications in which the user must remain in command of the system's operation. Indeed, in those difficult applications in which a reliable and rapid response is required, the value-driven methodology can provide a practical partnership between man and machine. A partnership that allows the man to retain flexible policy control over the system's behavior, while the automated system provides the required rapid response to real-time contingencies.

Application in Hierarchical Control Systems

Two significant properties of value-driven systems that make them particularly suitable for hierarchical applications are: the provision of a means for controlling the flow of information throughout the system, and the ability of each decision entity to continue operation in the absence of specific instructions from higher levels.

ORIGINAL PAGE IS
OF POOR QUALITY

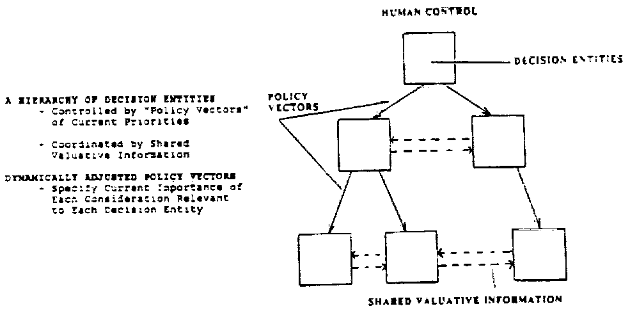

FIGURE 1. Fundamental Concept for Hierarchical Control

The basic design concept for the implementation of value-driven methodology within a computerized hierarchy is illustrated in Figure 1. Each activity is represented by a decision entity, which appears in the figure as a box. Each decision entity is guided by "policy vectors" that are defined at the next higher level in the appropriate control hierarchy. This incoming policy vector defines for the decision entity the current relative importance of the various policy considerations that are to be taken into account in the evaluation of alternatives. In order to coordinate the decisions at any level so as to most efficiently understand the state of the entire system, information is shared among the decision entities at that level, as shown by the dotted lines in the figure.

THE TAC BRAWLER AIR COMBAT SIMULATION

Introduction

DSA 's TAC BRAWLER model is a comprehensive simulation tool which provides a detailed representation of air-to-air combat engagements involving multiple flights of aircraft. Because of the importance of cooperative tactics and the critical role of human factors (such as surprise, confusion, situation awareness, and the ability to innovate tactical responses in unexpected situations) special emphasis has been placed on simulating these aspects of the engagement process.

To date, TAC BRAWLER has successfully reproduced the characteristics of engagements such as the ACEVAL-AIMVAL flight test series, and the manned simulator engagements made in conjunction with the AMRAAM OUE and has helped determine the characteristics that visual display systems for manned simulators must have if they are to provide realistic training in air-to-air combat. Presently, TAC BRAWLER is being extensively used in design studies for the Advanced Tactical Fighter program and for advanced avionics programs.

Modeling Pilot Behavior

The key factor to accurate modeling of air combat is the treatment of the human decision processes which drive the outcome of air-to-air engagements. DSA has developed a dual approach to the modeling of human decision processes which involves both value-driven decision-making and information-oriented decision architecture. This approach provides a practical solution to the problems involved in modeling multiple aircraft combat as in TAC BRAWLER, and includes:

- Explicit Model of Information Flow
  - Sensors
  - Communications
- Realistic Simulation of Decisions
  - Situation Assessment
  - Explicit Mental Model
  - Consideration of Alternatives
  - Based on Judgmental Values

The explicit simulation of the flow of information into and out of each pilot's personal mental model of the situation is key to the successful operation of TAC BRAWLER. Simulation of the pilot's decision process refers to value-driven decision making. Pilots use their mental models to perform situation assessment functions, generate sets of alternative courses of action, and select a particular action for implementation. The choice among alternatives is made on the basis of a judgmental evaluation of the situation that the pilot believes will result if a particular alternative is adopted.

Information-Oriented Architecture

Figure 2 shows the structure of the Information-oriented architecture of TAC BRAWLER. The central status arrays contain the true physical state of the simulation. Each simulated decision-maker has a personal mental status array which mirrors the central status. The imaging, however, is imperfect; a pilot will not know precisely where other aircraft are or exactly how fast they are moving. More important, aircraft and missiles of which he is unaware will be entirely absent from his mental model.

Information arrives in the mental model via sensor events and communications events which simulate visual, radar, missile warning receivers, radio communications, etc. TAC BRAWLER achieves

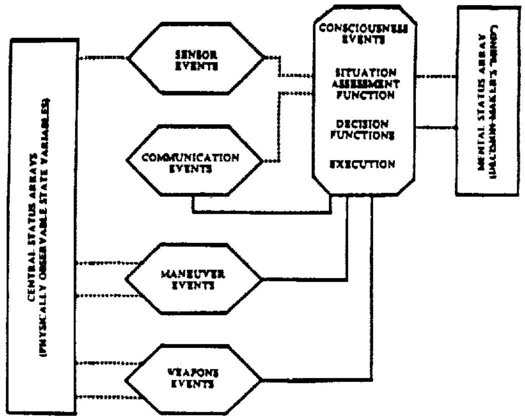

FIGURE 2. TAC BRAWLER Information-Oriented Architecture

a very high level of detail in the modeling of the physical characteristics of aircraft aerodynamics and systems.

Decisions by the pilot cause physical actions to occur either indirectly, via communications, or directly, through aircraft maneuver events and weapons employment events.

Hierarchy of Pilot Decisions

Figure 3 shows the hierarchy of decisions in TAC BRAWLER. The primary effect of high-level decisions is to control the lower level decisions by modifying their evaluation functions and by determining which lower level alternative actions will be considered.

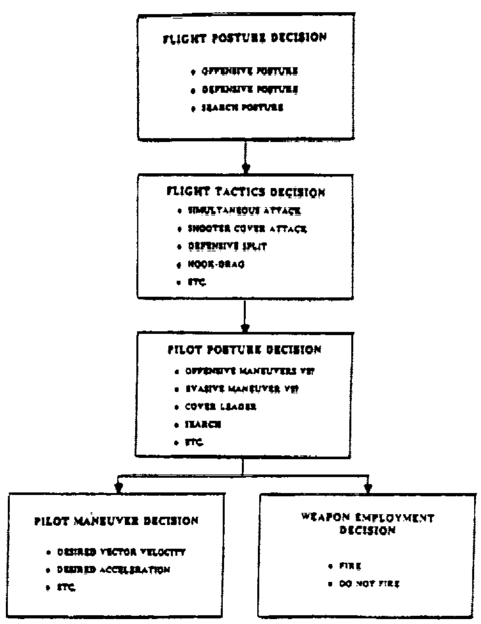

FIGURE 3. TAC BRAWLER Decision Hierarchy

The flight posture decision occurs at the highest level and determines the general course of action. It is made on the basis of a broad description of the situation, such as force ratios and engagement geometry, and also on the basis of user-supplied priorities. At the next level the flight leader determines the tactics that should be used to implement the flight posture. As a result of this decision a specific communication is sent to other flight members, informing them of the tactics. The effect of the message is to influence the values that other pilots use to score the alternative actions they consider. For instance, an order to attack a certain aircraft results in the subordinate perceiving that hostile as being more valuable; he "likes" to attack it. An interesting feature of this value-oriented representation of orders is that pilots can continue to make reasonable decisions in the absence of orders--they have an intrinsic, or default, set of values to use in their decision processes. Additional realism is present, because a pilot's perception of the situation may be very different from that of his flight leader. Since orders only influence a subordinate's mental model and do not force an action upon him, the subordinate exhibits a certain amount of common sense in his actions.

The pilot posture decision determines the general course of action for a pilot: to attack, to evade, to support the flight leader, etc. This decision controls the operation of aircraft systems such as radar and weapons selection, as well as deciding on the current weapons objective.

The maneuver decision and the weapon employment decisions occur at the lowest level. The alternatives considered by the maneuver decision are designed to do things like "get on the hostile's tail", "gain specific energy versus a hostile", "avoid the ground", etc. The weapon employment decision determines whether to fire a weapon at this moment.

Summary

The TAC BRAWLER simulation is the most mature and widely used application of value-driven hierarchical control methods. The model provides a synergy between value-driven logic to simulate pilot decision processes and an information-oriented hierarchical architecture as a means to represent the imperfect knowledge that a decision entity has of the true state of the situation. TAC BRAWLER has become a standard for high fidelity air-to-air engagement modeling within the Air Force, airframe manufacturing, and avionics communities. See reference 1 for a more complete description of the simulation.

ORIGINAL PAGE IS
OF POOR QUALITY

INTELLIGENT PLANNER FOR MULTIPLE AUTONOMOUS UNDERWATER VEHICLES

Introduction

DSA has developed a hierarchical value-driven methodology and demonstration prototype system to provide high-level decision and planning functions for the cooperating multiple autonomous underwater vehicles (AUVs) program.

If an AUV (or group of vehicles) is to operate on a truly autonomous basis, then the on-board computer system must be able to accomplish all of the high-level decision functions normally done by the human controller. These decision functions will involve some very complex activities--including even the replanning of the mission itself, in response to unexpected delays or events that may preclude the completion of the mission as originally anticipated. The on-board system must be able to develop and evaluate major revisions in the routing for the mission and it must be able to assess the need to make major changes in the basic mission plan--including even the omission of some or all of the originally planned tasks.

The AUV Operating Environment

The specific decision processes that are needed for the control of an autonomous vehicle depend fundamentally on the environment in which the vehicle is to operate--and on the types of threats and opportunities that are expected. The demonstration includes descriptions of: the natural environment for water depth profile and sonic transmission properties, and; individual vehicle descriptions of estimated position, energy state, and messages awaiting transmission. The postulated test scenarios accounted for potentially hostile objects such as:

- Mines and mine field areas
- Sonobuoys and sonobuoy areas
- Ships or submarines with active or passive sonar

The vehicles are concerned with the planning and execution of missions within this combined environment, in which the specified objectives may include:

- Specified area search and mapping
- Target identification and attack
- Detection avoidance
- The delivery to the operating base or host ship of any significant information from the mission.

Illustrative Tactical Options

The character of the decision-making functions is also dependent on the types of tactical options that are available to the vehicles. Individual vehicles can travel at low speeds (low noise), skirt areas of possible sonar detection, and hide in the bottom clutter. When travelling as a group in dangerous areas, these single vehicle tactics can be augmented by cooperative multivehicle tactics; for example:

- A "high low" tactic
- A "leader-follower"
- A "leap frog" tactic
- A separation with planned rendezvous.

Hierarchical Planning and Control Architecture

For the AUV project, DSA developed three levels of control which closely parallel the normal division of responsibility that is typical of a human command hierarchy: Mission Level, Group Level, and Vehicle Level.

MISSION LEVEL--THE GROUP PLANNERS

The mission level of the system receives a mission definition and basic values from the human Mission Planner. The tasks comprising the mission consist of such items as: search and destroy, reconnaissance, or map. These high level commands will then be broken into a series of subtasks to be performed by the group(s). It is the job of each Group Planner (at the mission level) to devise the actual group tactics and priorities to be employed in carrying out each of the mission subtasks.

GROUP LEVEL--THE VEHICLE PLANNERS

The Group Level decision entity takes the instructions from the Mission Level and turns them into detailed instructions for the Vehicle Levels below. There is one Vehicle Planner for each vehicle in the group, and an overall Planner Manager for the group as a whole. The Planner Manager will make decisions such as which vehicle is the leader and which is the follower in a "leader follower" tactic. A Planner Manager will also plan a route from the current position to the destination specified by the Mission Level.

The Group Level sends instructions to the Vehicle Level in the form of: a desired final location, a desired velocity, a desired depth, and a desired tactic.

VEHICLE LEVEL

The Vehicle Level decision entity takes the instruction from the Group Level and routes the vehicle along a detailed route at the chosen speed and depth. It is, in effect, the navigator of the system, plotting a course to the chosen destination (within the range of its sensors), moving the vehicle around small obstacles detected by the sensors.

The Vehicle Level decides on a optimal route for the vehicle based on its knowledge of the terrain, either through its preset information data base or through knowledge gained by its sensors.

Summary

DSA's prototype demonstration system implements the concepts described above, and shows how at each level the appropriate Planners make decisions based on a trade-off among the many factors that must be considered in mission planning, such as: urgency of mission, risk of detection, and safety of the vehicle.

SPACE STATION ELECTRIC POWER SCHEDULER

Introduction

DSA has recently completed a design concept for the automated control and scheduling of the Space Station Freedom (SSF) electric power system. The concept is different than that of the previous two examples in that it postulates the entire SSF environment as a free market economy where buyers and sellers of resources must bargain for their best options.

The electric power control system is only one of many automated subsystems that must be coordinated to provide a productive environment aboard SSF. To dispatch electric power to satisfy the demands of users without violating any resource constraints will require cooperative problem-solving among the subsystems, payloads, and the OMS to maximize productivity.

Planning and Scheduling

Reference 3 contains a detailed mathematical description of the value-driven approach. The approach may be described physically as a planning hierarchy of automated agents operating in a free-market environment:

- Computerized Resource Management Agents, one for *each* of the managers who are responsible for supplying resources such as electric power, thermal control, or life support for the Space Station.

- Computerized Resource Requesting Agents, one for *each* of the projects or activities that are major consumers of Space Station resources.

- A Free-Market Coordinating Agent responsible for managing and expediting the operation of the free-market as a resource allocation and scheduling mechanism.

The market coordinating agent initiates the process by postulating a trial set of time-dependent prices for each major resource. He then polls the negotiating agents for their individual responses. Each independent agent then responds with a specific plan for his own area of responsibility that would provide the best results for his area, given the postulated price structure.

The resource consuming agents respond first, by specifying how they would schedule their activities to maximize the profitability of their activities on the basis of the postulated price structure. Naturally, insofar as their benefits are not sensitive to the schedule they will schedule their activities to minimize their resource costs. But where the timing of an activity is important to its projected benefit or probability of success, the agent will select a scheduling alternative that maximizes the "profit" (i.e. the excess of the research benefit minus the resource costs). The combined scheduling responses of all of the resource consuming agents defines the overall schedule of consumer demand for each resource that would result from the postulated price structure.

The resource supplying agents respond next. Each agent responds with a specific plan for the delivery of his own resource (here, electric power) that would be most appropriate in the context of the foregoing schedule of consumer demand and the trial prices. If it is feasible to meet the schedule of demand, the resource supplying agent seeks to do so in the best and most efficient way, and he provides an estimate of the fair price that he would have to charge per unit of resource in each time period, in order to meet the specified demand.

Based on these combined responses, the market coordinator assesses the supply-demand and pricing relationships for each resource in each time period and adjusts his trial prices accordingly, either raising or lowering to bring supply in line with demand. He then again polls the agents for their responses.

The free-market coordination process is repeated until a price structure is found which shows a satisfactory relationship between the supply and demand for all resources in all time periods. When convergence is achieved, the free-market coordinator's trial prices will be equal (within an acceptable error tolerance) to the prices as estimated by the individual resource management agents. Therefore, in the final coordinated schedules the supply-demand relationships for each resource will properly reflect the resource supplier's expert estimate of both the costs and the risks of meeting an additional increment of demand. The estimation of these operating costs and the operating risks for various levels of demand (taking into account appropriate margins of safety) is, of course, one of the major responsibilities of the technical managers for the major Space Station resources.

Hierarchical Structure

In this approach the priorities, or values, of each project are defined in terms of a mathematical function whose parameters are divided into two major classes--those that must be controlled by the overall program manager, and those that can be controlled independently by the technical managers for the individual Space Station projects. The two classes of value parameters include the following types of considerations.

1. Parameters Controlled by the Overall Program Manager

    - The relative importance or priority assigned to each project

    - The estimated time urgency of each project

2. Parameters Controlled by the Technical Manager for Each Project

    - The relative technical value or utility of alternative research opportunities involving differences in the duration of the time allotted, in the specific orbits in which the time is allotted, in the specific time allotted within an orbit, and in the allotted envelopes of power and thermal consumption, etc.

    - The dependence of the technical value on other factors associated with scheduling dynamics, such as the time interval between allotted research opportunities, the cost of departing from a previously defined schedule, the importance of not interrupting a research activity once it is started, and so on.

The use of a formalized value structure allows the overall program manager to influence the relative scheduling priorities, without distorting the detailed scheduling preferences of the individual project managers; and it allows individual project managers to specify as broad or as narrow a range of scheduling preferences as they deem appropriate for their project. Perhaps most importantly, the approach provides an appropriate dynamic response to variations in the resource prices in which an increase in price can remove the low-priority or low-urgency projects from the scheduling competition, while at the same time allowing the most urgent or highest priority projects to continue at their required operating levels.

COMPARISON WITH EXPERT SYSTEMS

In its domain of application, value-driven control methods overlap the domain of the branch of Artificial Intelligence known as "expert systems." Expert systems focus on the actions to take at each step of a process. A system of rules is developed (often called situation-action rules) that determine what action to take, perhaps to consider another rule, at each step.

In contrast to this focus on the steps in the process, value-driven systems focus on the overall objective of the process. That is, value systems are goal-oriented, which has several advantages. For many systems it is desirable to be able to compare different policies. For example, in studying pilot behavior it is important to compare aggressive strategies--with a consequent increase in hostile kills--to more conservative strategies--with a consequent increase in friendly survivors. Or, more generally, it is often extremely important for a user to know what objective is being optimized. In rule-based systems this is often difficult to discern, particularly when new rules are added to a system.

Complex systems are frequently easier to model using value-driven techniques. For example, the Air Force model that was in use prior to TAC BRAWLER was rule-based; but it modeled only two aircraft. That is, it was a one-on-one model. Generalizing that model to represent multiple aircraft proved totally unwieldy; there were simply too many rules that had to be developed. The need for an exhaustive and consistent set of rules completely overwhelmed the situation. A value-driven system whose objectives are to a large extent independent of the number of aircraft proved much more efficient.

On the other hand, there are many situations that lend themselves very naturally to being represented by rule-based systems. These tend to have a linear or sequential nature. The medical diagnosis system MYCIN is of this type. For this type of system it makes a great deal of sense to proceed through a series of steps to arrive at a correct diagnosis, rather than attempting to optimize some difficult to define value function.

The significant difference between value-driven systems and expert systems seems to lie in the linear and sequential nature of the problems that lend themselves to efficient treatment by rule-based techniques. For, by contrast, problems that lend themselves to treatment by goal-oriented techniques tend to have a collective nature to them; no single (simple) path can be defined that leads directly to the solution.

COMPARISON WITH MATHEMATICAL PROGRAMMING

Value-driven theory is an outgrowth of a very general method for optimizing large systems that are characterized by non-linear objective functions, are defined over a discrete space, and are subject to inequality constraints. The method known as Generalized Lagrange Multiplier (GLM) Theory, and described in reference 4, is a generalization of the classical Lagrange method. As such, value-driven systems are a form of mathematical programming. However, value-driven systems, including even those that are specifically designed as optimizing systems, are typically quite different from the more widely known techniques in both their design philosophy and their operating characteristics.

In the application of most optimizing methods, it is usually necessary to make simplifying assumptions in the problem structure, to recast it in a standardized form that is compatible with an established procedure. That is, it is standard practice to sacrifice accuracy in the representation of the problem, in order to provide a rigorous and accurate mathematical optimization.

In a typical value-driven system, the design priorities are reversed. That is, the achievement of a precise mathematical optimum is rarely a driving objective in the problem formulation. Rigorous optimization tends to be sacrificed when necessary, to provide a more accurate representation of the problem including:

- A more accurate representation of the physical problem and its associated constraints, and

- A satisfactory representation of all of the valuative considerations that realistically must be taken into account in order to select an appropriate real-world course of action.

The focus on problem representation is aided immensely by some features of the GLM method itself. The lack of restriction on the objective (value) function--it can be almost any function imaginable--gives the developer great flexibility in realistically representing the complex objectives that must be reflected in real-world systems.

SUMMARY

The distinctive features of the value-driven methodology: focus on system-wide goals, responsiveness to command priorities, decomposition to decision entities, and applicability to large-scale systems, makes it possible to develop automated systems that can cope effectively with complex planning and control tasks that can be addressed only through a hierarchical decision process.

The three program described illustrate the three stages of system development for hierarchical control systems. A conceptual design for the automated scheduling of the Space Station Freedom electric power system. A prototype demonstration system for cooperating multiple autonomous underwater vehicles. A mature widely used simulation of multiple aircraft air-to-air combat.

REFERENCES

1. Kerchner, R. M., "The TAC BRAWLER Air Combat Simulation Management Summary Manual (Rev 5.0)," DSA Report #907 for HQ USAFCSA/SAGF, Contract F49642-87-C-0022, Decision-Science Applications, Inc. Arlington, VA, June 1988.

2. Krupp, Joseph C., "Providing Intelligent Planning for Multiple Autonomous Underwater Vehicles," DSA Report # 839 for National Bureau of Standards, Contract 50SBN6C4130, Decision-Science Applications, Inc. Arlington, VA, September 1987.

3. Pugh, George E., "A Value-Driven Electric Power Scheduler," DSA Report #977 for NASA Lewis Research Center, Decision-Science Applications, Inc. Arlington, VA, June 1989.

4. Evert, Hugh III, "Generalized Lagrange Multiplier Method for Solving Problems of Optimal Allocation of Resources," Journal of the Operations Research Society of America, Volume III, No. 3, May-June 1963.